



(11) **EP 3 291 103 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
06.11.2019 Bulletin 2019/45

(51) Int Cl.:
G06F 12/02 ^(2006.01) **G06F 12/1009** ^(2016.01)
G06F 16/23 ^(2019.01)

(21) Application number: **16186857.5**

(22) Date of filing: **01.09.2016**

(54) **SYSTEM AND METHOD FOR CREATING A SNAPSHOT OF A SUBSET OF A DATABASE**

SYSTEM UND VERFAHREN ZUR ERZEUGUNG EINER MOMENTAUFNAHME EINER TEILMENGE EINER DATENBANK

SYSTEME ET PROCEDE DE CREATION DE COPIES INSTANTANEEES D'UN SOUS-ENSEMBLE DE BASE DE DONNEES

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

(43) Date of publication of application:
07.03.2018 Bulletin 2018/10

(73) Proprietor: **Huawei Technologies Co., Ltd.**
Longgang District
Shenzhen, Guangdong 518129 (CN)

(72) Inventors:
• **GOIKHMAN, Shay**
80992 Munich (DE)
• **ILIOPOULOS, Antonios**
80992 Munich (DE)

• **LEVY, Eliezer**
80992 Munich (DE)

(74) Representative: **Kreuz, Georg Maria**
Huawei Technologies Duesseldorf GmbH
Riesstraße 25
80992 München (DE)

(56) References cited:
US-A1- 2013 073 513

• **KEVIN ELPHINSTONE ET AL: "Supporting Persistent Object Systems in a Single Address Space", PROCEEDINGS OF THE 7TH INTERNATIONAL WORKSHOP ON PERSISTENT OBJECT SYSTEMS, 1997, XP055337844,**

EP 3 291 103 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

DescriptionTECHNICAL FIELD

5 **[0001]** The present invention relates to a system for creating a snapshot and to a method for creating a snapshot. The present invention also relates to a computer-readable storage medium storing program code, the program code comprising instructions for carrying out such a method.

BACKGROUND

10 **[0002]** Creating a snapshot of a database is required for example when creating a backup of the database or when a query is to be executed on a consistent state of the database. In many cases, it is desired that the snapshot only comprises a subset of the data of the database.

15 **[0003]** As an example, a hybrid OLTP-OLAP database can utilize the *fork()* system call for creating a child process from a parent process serving as a consistent snapshot for the OLAP part for the duration of a long-running query execution. The snapshot can be maintained dynamically by means of the CoW (copy-on-write) optimization: the fork system call maps the shared page frames in both processes, the parent process and the child process, as read-only, and only when a page is modified by the parent process, the kernel allocates a new page frame, copies the page's data to the new frame and modifies the data in the new frame.

20 **[0004]** The same mechanism can be utilized for checkpointing for hybrid and other data stores. The checkpointing flushes a consistent database snapshot in the child process to persistent storage. This child process can also be a long-running one, as it depends on the size of the snapshot and the speed of the I/O.

[0005] In the above prior art methods, the computational effort for creating the snapshot is significant.

25 **[0006]** Document "Supporting Persistent Object Systems in a Single Address Space", in Proceedings of the 7th International Workshop on Persistent Object Systems, 1997, discloses that stability can be implemented in the Mungi SASOS (Single address space systems) without incurring overhead in excess of the inherent cost of shadow-paging. This paper is based on the introduction of aliasing into the SASOS model and makes heavy use of user-level page fault handlers to allow implementation outside the kernel. This paper also shows how the demands of database systems for control over page residency and physical I/O can be accommodated. An approach to user-level implementation of distributed shared memory (DSM) coherency models is outlined.

30 **[0007]** Document US 2013/073513 A1 discloses a method of maintaining a hybrid OLTP and OLAP database, the method comprising: executing one or more OLTP transactions; creating a virtual memory snapshot; and executing one or more OLAP queries using the virtual memory snapshot. Preferably, the method further comprises replicating a virtual memory page on which a data object is stored in response to an update to the data object, whereby the updated data object is accessible for OLTP transactions, while the non-updated data object remains accessible for OLAP queries. Accordingly, the present invention provides a hybrid system that can handle both OLTP and OLAP simultaneously by using hardware-assisted replication mechanisms to maintain consistent snapshots of the transactional data.

SUMMARY OF THE INVENTION

40 **[0008]** The objective of the present invention is to provide a system for creating a snapshot and a method for creating a snapshot, wherein the system for creating a snapshot and the method for creating a snapshot overcome one or more of the above-mentioned problems of the prior art.

45 **[0009]** A first aspect of the invention provides a system for creating a snapshot of a subset of a database that is stored as a plurality of segments in a process address space, the system comprising:

- a determining unit configured to determine from the plurality of segments one or more source segments of the snapshot, and
- a virtual memory creation unit configured to allocate for the snapshot one or more snapshot segments within the process address space, wherein the snapshot segments are mapped to same physical memory addresses as the source segments and wherein the source segments and/or the snapshot segments are set to copy-on-write; characterized in,

55 the system comprises further a checkpoint module, the checkpoint module being configured to determine the one or more source segments by determining which segments have been changed in a previous checkpoint interval, wherein the one or more snapshot segments comprise at least the segments which have been changed in the previous checkpoint interval.

[0010] The system of the first aspect departs from the legacy semantics traditionally associated with the forking

mechanism, by introducing a copy-on-write mechanism that operates within a single address space, rather than creating a separate address space. This approach enables the application to have fine-grained, lightweight and ad-hoc shadowing (versioning) of memory regions, without the overheads associated with setting up any additional system/application functionality related to a separate address space.

5 **[0011]** The source segments of a snapshot are those segments in the database for which the snapshot is taken, that is the segments that are duplicated. It is understood that in different application contexts there may be different ways of determining the source segments of a snapshot. For example, source segments of a snapshot can be selected from consecutive segments in the process address space. In this regard, the system according to the first aspect is not limited to a specific way of determining the snapshot segments.

10 **[0012]** In a first implementation of the system for creating a snapshot according to the first aspect, the virtual memory creation unit is implemented in the kernel of an operating system of the system.

[0013] This represents a particularly efficient way of implementing the allocating the snapshot segments. For example, the kernel can be the Linux kernel.

15 **[0014]** In a second implementation of the system for creating a snapshot according to the first aspect as such or according to the first implementation of the first aspect, the system is configured to, on an access to a segment which maps to a persistent storage:

- allocate one or more page frames,
- read the segment from the persistent storage to the page frames, and
- 20 - map the segment and a corresponding snapshot segment to the page frames.

[0015] This makes it possible that the system can be applied also for segments which are stored on a persistent storage. The persistent storage can be internal or external to the system. The persistent storage can be for example a hard drive and/or a solid state drive (SSD).

25 **[0016]** In a third implementation of the system for creating a snapshot according to the first aspect as such or according to any of the preceding implementations of the first aspect, the system further comprises a memory management unit that is configured to induce a page-fault when a source segment and/or a snapshot segment is accessed.

30 **[0017]** This represents a particularly efficient way of implementing the Copy-on-write. For example, if the snapshot segments are set to Copy-on-write, the memory management unit can be configured to induce a page fault when the consumer unit intends to write to one or more of the snapshot segments.

[0018] Preferably, the system further comprises a consumer unit, in particular a consumer thread that is configured to access the snapshot segments.

[0019] This allows a logical separation between the consumer unit that accesses the snapshot segments and the virtual memory creation unit that allocates the snapshot segments.

35 **[0020]** In a fourth possible implementation of the system for creating a snapshot according to the third implementation of the first aspect, the operating system is further configured to perform copy-on-write, COW, in response to the page fault, wherein the COW performing comprises allocating at least one page frame, copying the content of one or more memory pages referenced by the accessed snapshot segments to new memory pages, and setting the corresponding source segments to map to the new memory pages. Thus, the operating system can then handle the page fault by copying the content of one or more memory pages referenced by the accessed snapshot segments to new memory pages, and set the corresponding source segments to map to the new memory pages.

40 **[0021]** In a fifth implementation of the system for creating a snapshot according to the first aspect as such or according to any of the preceding implementations of the first aspect, the system comprises a compiling unit that is configured to compile a physical execution plan of a database query, in particular an OLAP query, wherein preferably the compiled physical execution plan comprises an unmap command for a snapshot of a segment at a position in the physical execution plan when the snapshot of the segment is no longer accessed during a remaining execution time of the physical execution plan.

45 **[0022]** In a sixth implementation the system for creating a snapshot according to the first aspect as such or according to any of the preceding implementations of the first aspect, the virtual memory creation unit is further configured to provide a single page table that corresponds to the process address space, wherein the page table comprise one or more page table entries for the source segments and one or more page table entries for the snapshot segments, and the page table entries of the snapshot segments are mapped to same physical memory addresses as the page table entries of the source segments to be snapshot.

50 **[0023]** In a seventh implementation of the system for creating a snapshot according to the first aspect as such or according to any of the preceding implementations of the first aspect, the virtual memory creation unit is further configured to set a read-only flag of the source segments, wherein the read-only flag indicates that the source segments are set to copy-on-write.

[0024] In an eight implementation of the system for creating a snapshot according to the first aspect as such or

according to any of the preceding implementations of the first aspect, the virtual memory creation unit is further configured to set a read-only flag of the snapshot segments, wherein the read-only flag indicates that the snapshot segments are set to copy-on-write.

[0025] The system of this implementation has the advantage that it provides a particularly efficient implementation of executing database query, where a snapshot of a current state of the database is typically required.

[0026] The physical execution plan can be compiled, from, e.g., an SQL representation to any physical plans representation: a DAG of predefined functions calls, or an LLVM intermediate code that can be executed by an interpreter, or the native machine code. For example, one can produce a function in native machine code that takes snapshot segments as arguments, such that executing this function produces the query result. Later, this function can be passed to some thread, along with the actual arguments and the thread executes it.

[0027] A second aspect of the invention refers to a method for creating a snapshot of a subset of a database that is stored as a plurality of segments in a process address space, the method comprising:

- determining from the plurality of segments one or more source segments of the snapshot, and
- allocating for the snapshot one or more snapshot segments within the process address space, wherein the snapshot segments are mapped to same physical addresses as the source segments and wherein the source segments and/or the snapshot segments are set to copy-on-write;

wherein determining the one or more source segments comprises a step of determining which segments have been changed in a previous checkpoint interval, wherein the one or more snapshot segments comprise at least the segments which have been changed in the previous checkpoint interval.

[0028] The methods according to the second aspect of the invention can be performed by the system according to the first aspect of the invention. Further features or implementations of the method according to the second aspect of the invention can perform the functionality of the system according to the first aspect of the invention and its different implementation forms.

[0029] In a first implementation of the method for creating a snapshot of the second aspect, the method further comprises, on an access to a segment which maps to a persistent storage:

- allocating one or more page frames,
- reading the segment from the persistent storage to the page frames, and
- mapping the segment and the corresponding snapshot segment to the page frames.

[0030] In a second implementation of the method for creating a snapshot of the second aspect as such or according to the first implementation of the second aspect, determining the one or more source segments comprises a step of determining which segments have been changed in a previous checkpoint interval.

[0031] The method of this implementation is particularly useful for implementing a checkpointing system which e.g. creates checkpoints in regular time intervals. By determining which segments have been changed in a previous checkpoint interval, the method can determine the source segments that are required for an incremental checkpoint. This has the advantage that the size of the checkpoint is kept as low as possible while still ensuring that a complete checkpoint is created.

[0032] In a third implementation of the method for creating a snapshot of the second aspect as such or according to any of the preceding implementations of the second aspect, the method is a method for maintaining a hybrid online transaction processing, OLTP, and online analytical processing, OLAP, database system and wherein the method further comprises receiving and executing an OLAP query.

[0033] In a fourth implementation of the method for creating a snapshot of the second aspect as such or according to any of the preceding implementations of the second aspect, determining the one or more source segments of the snapshot comprises a step of creating a physical execution plan, wherein in particular the source segments are referenced in the physical execution plan.

[0034] The physical execution plan comprises information about which segments are required for carrying out the database query. Thus, the method can precisely determine the snapshot segments.

[0035] In a fifth implementation of the method for creating a snapshot of the second aspect as such or according to any of the preceding implementations of the second aspect, the method further comprises:

- determining when a snapshot segment is no longer accessed by a consumer thread, and
- unmapping the snapshot segment.

[0036] Unmapping the snapshot segment can refer to de-allocating snapshot segments.

[0037] In a sixth implementation of the method for creating a snapshot of the second aspect as such or according to

any of the preceding implementations of the second aspect, the method further comprises:

- setting a read-only flag of the source segments and/or the snapshot segments,
- inducing a page-fault when one of the source segments is accessed and/or when one of the snapshot segments is accessed, and
- performing copy-on-write in response to the page fault.

[0038] Herein, the access refers to a write access.

[0039] The page fault can be dealt with for example by the operating system.

[0040] A third aspect of the invention refers to a computer-readable storage medium storing program code, the program code comprising instructions for carrying out the method of the second aspect or one of the implementations of the second aspect.

BRIEF DESCRIPTION OF THE DRAWINGS

[0041] To illustrate the technical features of embodiments of the present invention more clearly, the accompanying drawings provided for describing the embodiments are introduced briefly in the following. The accompanying drawings in the following description are merely some embodiments of the present invention, but modifications on these embodiments are possible without departing from the scope of the present invention as defined in the claims.

FIG. 1 is a block diagram illustrating a system for creating a snapshot in accordance with an embodiment of the present invention;

FIG. 2 is a flow chart illustrating a method for creating a snapshot in accordance with an embodiment of the present invention;

FIG. 3 illustrates segments of a page table and the physical memory pages that the segments map to, in accordance with an embodiment of the present invention;

FIG. 4 shows a schematic illustration of an example of a physical execution plan of two joins of three tables of a database, in accordance with an embodiment of the present invention;

FIG. 5 shows a schematic illustration of a method of determining segments that are accessed by a snapshot process of a checkpointing method in accordance with an embodiment of the present invention; and

FIG. 6 shows a schematic illustration of a system for creating a snapshot of a subset of a database in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION OF THE EMBODIMENTS

[0042] FIG. 1 shows a system 100 for creating a snapshot of a subset of a database that is stored as a plurality of segments in an address space. A segment comprises one or more memory pages.

[0043] The system 100 comprises a determining unit 110 and a virtual memory creation unit 120.

[0044] The determining unit 110 is configured to determine from the plurality of segments one or more source segments of the snapshot.

[0045] The virtual memory creation unit 120 is configured to allocate for the snapshot one or more snapshot segments within the single or same address space. Therein, the snapshot segments are mapped to same physical memory addresses as the source segments and wherein the source segments and/or the snapshot segments are set to copy-on-write.

[0046] In a traditional scenario, an operating system forking mechanism can operate across two or more distinct and separate address spaces in virtual memory, semantically forming a parent - child relationship. During the fork operation, the operating system kernel creates a new address space and a new page table for the child process. Further, the operating system kernel duplicates the parent's page table hierarchy and page table entries to the newly created page table that belongs to the child.

[0047] The aforementioned forking mechanism is further leveraged in an unconventional manner at settings where an application requires a consistent memory snapshot that may be provided lazily by means of the copy-on-write semantics. In such cases, leveraging the traditional forking mechanism of an operating system bears several performance disadvantages, as the forking mechanism was not designed to serve the particular case of providing memory snapshots

in mind, a case which admits several performance improvements over the common case. In particular, the traditional forking mechanism is pessimistic in nature; that is, the mechanism creates a new virtual address space (representing the child process), and duplicates the entire page table hierarchy via copying all page table entries from the parent, and further marking all of them with read-only protection bits. This incurs significant overheads at runtime for an application that wishes to leverage the mechanism only for the purposes of memory snapshots, but does not depend on having a parent-child relationship (i.e. does not require the creation of a new process, and its accompanying overheads). As such several applications benefit from having a more fine-grained control over the areas of interest to be snapshot thereby reducing the duplication of the entire page table, as well as from avoiding the creation of another separate process (i.e. child process) for each snapshot, by having the snapshot residing on the same address space as the process that requests it.

[0048] In particular, the system 100 can benefit from accessing the snapshot data directly in its own address space, instead of accessing the data across a separate address space over an inter-process communication channel that incurs extra overheads (in the case of the traditional parent-child forking mechanism). This saves address space, as only a single address space is required.

[0049] The system 100 is based on the concept of a single-address-space copy-on-write snapshot mechanism, able to accommodate the requirements of an application that requires consistent memory snapshots, while avoiding the aforementioned overheads of leveraging the traditional forking mechanism. The system 100 can provide an interface to applications (not shown in FIG. 1) from the operating system and the interface allows for the applications to designate any portions of their address space as memory ranges to be "shadowed" by employing a copy-on-write mechanism within the same address space. The granularity of the ranges is in quantities of memory pages (as supported by the underlying architecture).

[0050] Preferably, the system 100 can be implemented in an operating system kernel. In particular, operating system kernel virtual memory management mechanisms can be enhanced in order to provide the means to supply a copy-on-write based memory snapshot on the address space of the requesting process. The operating system kernel reserves an unused range of the virtual memory address space of the process, equivalent in size to that of the memory region that the process requested to be snapshot. The operating system kernel further creates new page table entries within the page table of the requesting process, in order to accommodate the region which is being snapshot. The exact number of the page table entries that is created is identical to the number of page table entries that cover the memory region to be snapshot. The newly created page table entries cover the virtual address space memory range reserved for the snapshot, and they are created to point to the same underlying memory pages as the page table entries that cover the memory region to be snapshot. The page table entries of both the original memory range and those of its snapshot are all further modified in order to mark their protection bits as read-only. At this point, the mechanism may return control to the process. The process may subsequently access both the original memory region and the shadowed memory region within the process' own address space. The original memory region may correspond to the source segments, and the shadowed memory region may correspond to the snapshot segments. The copy-on-write semantics ensure that write access to either the original or the shadowed memory range will induce a page-fault via the underlying hardware platform's MMU. The page-fault will be handled by the operating system kernel accordingly, which is able to discern that the fault concerns a page that is under copy-on-write semantics (much like the traditional forking mechanism case). The kernel then proceeds to allocate a new page frame, make a copy of the original page onto the newly allocated page frame, and modify the corresponding page table entry of the page that corresponds to the original memory region to point to the newly allocated and copied page.

[0051] Finally, the kernel modifies the protection bits of both page table entries (the one that belongs to the original memory region and the corresponding one that belongs to the shadowed memory region), and sets them to read-write, after which the page-fault handling concludes, and control may return to the process which can continue to access its original as well as its snapshot memory transparently.

[0052] In one embodiment, the system is implemented within the Linux operating system kernel. In particular, the virtual memory management subsystem is extended, and new primitives are exposed by means of introducing a set of new system calls. The main system call allows an application process to request a memory range of interest to be snapshot. The memory range is communicated to the kernel via the system call by providing a starting virtual address (that belongs to the process' own virtual address space), along with a second parameter which denotes the length of the memory region to be snapshot. Internally, the Linux kernel, based on the memory region address and length, detects the corresponding VMAs (virtual memory areas), which signify the reserved virtual address space memory regions that the kernel has allocated for that particular memory region. The kernel then proceeds to locate an unreserved region in the process's virtual address space which should be as large as the length of the snapshot. This operation is assumed to always succeed in a 64-bit address space. Subsequently, the kernel modifies the page table of the process and creates page table entries to cover the newly reserved virtual address space region to accommodate the snapshot. The page table entries are made to point to the same physical page frames as the page table entries of the original memory region that they are requested to shadow. Both the page table entries of the original memory region, as well as the

shadow memory region are further modified in order to set their protection bits to read-only, in order to enforce the copy-on-write semantics.

[0053] The system 100 can make use of virtual memory primitives that enable shadowing (versioning) a virtual space range inside a process while setting it copy-on-write. Moreover, it can redefine the handling of the demand-paging associated with the shadowing, such that no additional page-fault and I/O is needed in the case that the shadowed page has been demand-paged and then has been reclaimed, and later the shadowing page is accessed.

[0054] These primitives can enable consistent snapshots of selected regions of the database within the same process. Thus, the same pool of threads, allocated once, can service both OLTP and OLAP queries and the checkpoints. An application-level monitor can load balance the number of threads per task and for a global utility metric. The meta-data control information can be natively shared among all the application threads or snapshotted as well.

[0055] In particular, the following new virtual memory primitives can be defined:

void *shadow_vma = create_vma_shadow(int nr_vmas, int RW_flag, void *vma, ...)

[0056] *nr_vmas* designates the number of source *vma* segments. *create_vma_shadow* allocates a new virtual address area in the process's address space, binding it to the physical pages of the one or more source *vma(s)* and setting all the source *vma(s)* copy-on-write if the *RW_flag* is 0. Otherwise, if the *RW_flag* is 1, the returned *shadow_vma* is set copy-on-write. If any of the *vma(s)* pages is mapped to a file and therefore there is no page table entry (PTE) as yet for a page at the time of the *create_vma_shadow* call, on the first page access a frame will be allocated and its contents read from the file, and the corresponding PTE will be placed both in the *vma* and in the *shadow_vma* with the read-only flag according to the call semantics. Subsequently, the memory access will be processed as defined above.

[0057] **void reclaim_vma_shadow(void *shadow_vma)** reclaims the virtual space occupied by the shadow *vma* and frees any orphan pages associated with it.

[0058] As an additional advantage of this method, it is able to provide virtually contiguous memory area snapshot for potentially virtually discontinuous and distinct source virtual memory areas. This may be leveraged to create ad-hoc "views" of data with low overhead (by avoiding memory copying). It may furthermore be used to create ad-hoc virtually contiguous areas of data that are to be scanned serially in memory, thereby enabling the memory scans to run more efficiently, due to the introduced contiguity of the virtual addresses which can be fully predicted and prefetch.

[0059] The enhancement of the virtual memory subsystem enables taking a consistent snapshot of selected data within the same process (i.e., within a single virtual address space). We utilize this feature to devise a hybrid OLTP-OLAP database, a checkpointing data store, or any application that requires taking a consistent snapshot, to run within the same process.

[0060] FIG. 2 shows a method 200 for creating a snapshot of a subset of a database that is stored as a plurality of segments in an address space.

[0061] The method comprises a first step of determining 210 from the plurality of segments one or more source segments of the snapshot.

[0062] The method comprises a second step of allocating 220 for the snapshot one or more snapshot segments within the address space, wherein the snapshot segments are mapped to same physical addresses as the source segments and wherein the source segments and/or the snapshot segments are set to copy-on-write.

[0063] FIG. 3 shows how source segments 312, 314 of a database in a page table 310 map to physical pages 320, 322. Segment A 312 maps to a first set of (three) physical pages 320. Segment B 314 maps to a second set of (three) physical pages 322. Furthermore, the page table 310 comprises entries for a shadow virtual memory area 330. This shadow virtual memory area 330 comprises a snapshot 332 of segment A and a snapshot 334 of segment B. It is noted that one or more entries of the page table 310 refer to one or more VMAs of one or more segments. In particular, the relative position of the PTE defines its virtual address. For example, as the page table is an array of PTEs each taking 4k of virtual address space, the fifth PTE would then map virtual addresses 20k to 24k-1 to some page frame.

[0064] With other words, FIG. 3 shows a page table 310 comprising page table entries 312, 314 of an original virtual memory area and page table entries 332, 334 of a shadow virtual memory area 330. The original virtual memory area comprises segment A and segment B. The shadow virtual memory area 330 comprises a snapshot of segment A (i.e. snapshot segment A') and a snapshot of segment B (i.e. snapshot segment B'). The page table entries 332 of a snapshot of segment A are mapped to the same physical page frames 320 as the page table entries 312 of the segment A. The page table entries 334 of a snapshot of segment B are mapped to the same physical page frames 322 as the page table entries 314 of the segment B.

[0065] In an embodiment of the invention, the system comprises a compiling unit that is configured to compile a physical execution plan of a database query. FIG. 4 illustrates an example of a physical execution plan of two joins of three tables R, S, and U. The physical execution plan includes the designation of the segments that the query operates on.

[0066] In particular, the physical execution plan includes a step 416 of performing a TableScan of the segments of

the table R, and a step 418 of performing a TableScan of the segments of the table S. Subsequently, in step 412, a hash-join is performed on a materialized result of step 416 and 418, and in step 414 a TableScan is performed on segments of table U. Finally, on the results of the two-pass hash-join of step 412, in step 410 a TableScan is performed on the materialized result of steps 414 and 412.

5 **[0067]** The compiling unit accepts OLAP SQL query statements, which it parses, analyses and rewrites, transforming the query parse tree using the database metadata, into a physical execution plan, e.g., a tree of operator function calls and their arguments.

10 **[0068]** The checkpointer module runs an incremental checkpointing algorithm. The time is divided into checkpoint intervals, when at each end of an interval a snapshot constituting the memory segments that were modified throughout the interval is taken by requesting the snapshot manager to set the segments as inheritable and call fork(). The checkpoint process outputs these segments to the disk. The process returns success status if all the segments were persisted successfully. Otherwise, the checkpointer module will accumulate the segments that failed to be persisted in previous intervals with the ones modified throughout the current interval in the same snapshot to be persisted at the end of the current interval.

15 **[0069]** FIG. 5 provides an illustration of a method for checkpointing across several checkpoint intervals. A plurality of four objects 531, 532, 533 and 534 are tracked over four time intervals. The first, second and third time interval are indicated with reference numbers 511, 512, and 513 in FIG. 5. The first time interval 511 precedes the first checkpoint 521. The second and third time intervals 512, 513 are defined as the intervals between first, second and third checkpoint 521, 522, 523. The fourth time interval comes after the third checkpoint 523. The first and fourth objects 531, 534 are changed (indicated with reference numbers 541, 542) in the first interval 511. The fourth object is changed 543 again in the second interval 512. In the third interval 513, the second and third object 532, 533 are changed 544, 545.

20 **[0070]** In the first time interval 511 (that ends at the time of the first checkpoint 521), the change sets consist of the first and fourth object 531, 534, i.e., the objects that were modified before the first checkpoint 521. At the time of the first checkpoint, i.e., when the second time interval is started the objects are marked as changed relatively to the first checkpoint. The change set consisting of the first and fourth object 531, 534 is persisted successfully during the second time interval.

25 **[0071]** On the outset of the third time interval the change set consisting of the fourth object is detected, i.e., it is detected that a fourth object has been modified between the first and second checkpoint 521, 522. It is attempted to flush the second change set to disk. That results (for some unknown reason) in a disk write failure. In the fourth time interval 514, the computed change set consists of the second, third and fourth object 532, 533, 534. This change set results from the union of the fourth object 534 that failed to persist in the third time interval 513 and the change set consisting of the second and third object 532, 533 that is detected as changed in the third time interval 513.

30 **[0072]** In general, applying the 'separation of concerns' concept, the method of snapshot capture needs not being of the checkpointing method's concern, as long as a consistent immutable replica of the state of the objects is created.

35 **[0073]** At the outset of a new checkpoint, the status variables of the objects that have changed relatively to the last checkpoint need to be retained, together with the marking of all the objects as unchanged for the duration of the forthcoming checkpoint, all these in the presence of multiple parallel updates.

40 **[0074]** In a preferred embodiment of the invention, a system for creating a snapshot of a subset of a database comprises a snapshot manager which is responsible for managing snapshot both for an OLAP query planner module and for a check-pointing module.

[0075] This embodiment assumes that the mentioned modules are part of the same hybrid OLTP-OLAP database process. The OLAP query planner module, upon receiving the query, determines which memory regions of contiguous pages, called segments, are required in the query executor. Similarly, the checkpointer module determines the segments needed in the incremental checkpointing algorithm executed in the checkpointer module.

45 **[0076]** The OLAP query planner module accepts OLAP SQL query statements, which it parses, analyses and rewrites, transforming the query parse tree using the database metadata, into a physical execution plan, e.g., a tree of operator function calls and their arguments. The physical plan includes the designation of the segments the query operates on, see e.g. FIG. 4. Before running the query executor, the query planner requests the snapshot manager to shadow the segments referenced in the physical plan. The shadowed regions constitute a consistent snapshot which is read by the query executor, while at the same time the OLTP part can modify the segments.

50 **[0077]** The checkpointer module runs an incremental checkpointing algorithm. The time is divided into checkpoint intervals, when at each end of an interval a snapshot constituting the memory segments that were modified throughout the interval is taken by requesting the snapshot manager to shadow the segments. The checkpointer module using one or more threads outputs these segments to the disk. The checkpointer returns success status if all the segments were persisted successfully. Otherwise, the checkpointer module will accumulate the segments that failed to be persisted in previous intervals with the ones modified throughout the current interval in the same snapshot to be persisted at the end of the current interval.

55 **[0078]** FIG. 6 illustrates a system 600 comprising a query planner module 602, a checkpointing module 604 and a

snapshot manager 620. The query planner module 602 comprises a first determining unit 603 for determining which segments will be accessed by a query executor process 630. The query planner module 602 is configured to pass these segments 610 that are accessed by the query executor process 630 to the snapshot manager 620. Similarly, the checkpointing module 604 comprises a second determining unit 605 for determining which segments will be accessed by a checkpointing process 632. The checkpointing module 604 is configured to pass segments 612 to the snapshot manager 620.

[0079] The query planner module 602 is configured to determine the segments appearing in its physical plan. The checkpointing module 604 tracks the segments changed in the checkpointing interval. Each of them in its turn passes its segments to the snapshot manager 620.

[0080] The snapshot manager 620 comprises a virtual memory creation unit 622. The virtual memory creation unit 622 is configured to allocate for the segments 610 from the query planner module 602 and for the segments 612 from the checkpointing module 604 a plurality of snapshot segments within a current address space. As part of a vma-shadow operation, the snapshot segments are made available to a query executor module 630 and a checkpointer module 632.

[0081] In embodiments of the invention, the snapshot manager 620 can provide an interface with a function for initializing the creation of a new snapshot process, a function for adding segments and/or a function for finalizing the creation of the snapshot process.

[0082] Embodiments of the invention can have the following advantages:

Some embodiments can enable low-overhead snapshotting (versioning) of selected memory regions into the same address space. As such, and moreover, it can be used as an efficient means for process creation in a single-address-space OS.

[0083] Some embodiments can avoid complexity, overhead and inefficiency in managing and sharing information and resources between modules requiring snapshot and comprising the same application. E.g., the method enables the initialization of the OLTP part and the OLAP part and the Checkpointer parts to happen only once, whereas in the fork-based method the initialization of OLAP part occurs once per each child forked. The method can avoid redundant system resource over-committing and system processing, e.g., the same demand-paging happening multiple-times.

[0084] Some embodiments can enable efficient resource and information sharing among all the application modules requiring snapshot. As such, the method can enable adaptive load balancing among all the parts of an application, e.g., hybrid OLTP-OLAP database, in order to achieve any globally defined utility metric.

[0085] Some embodiments can enable the hybrid OLTP-OLAP database to provide more frequent and fresher snapshots for the analytical (OLAP) queries, utilizing lightweight snapshotting functionality. The low overhead of the mechanism can enable the hybrid database to dynamically provide ad-hoc snapshots (per OLAP query) rather than period snapshots.

[0086] Some embodiments enable creating virtually contiguous memory regions out of potentially virtual discontinuous page-aligned memory regions. Those regions may be leveraged by the database to provide e.g. ad-hoc data views with very low overhead, avoiding memory copying and/or more efficient memory scans.

[0087] The foregoing descriptions are only implementation manners of the present invention, the scope of the present invention is not limited to this. Any variations or replacements can be easily made through person skilled in the art. Therefore, the protection scope of the present invention should be subject to the protection scope of the attached claims.

Claims

1. A system (100) for creating a snapshot of a subset of a database that is stored as a plurality of segments (312, 314; 531-534) in a process address space, the system comprising:

- a determining unit (110) configured to determine from the plurality of segments one or more source segments (312, 314; 531-534) of the snapshot, and

- a virtual memory creation unit (120) configured to allocate for the snapshot one or more snapshot segments (332, 334) within the process address space, wherein the snapshot segments are mapped to same physical memory addresses as the source segments and wherein the source segments and/or the snapshot segments are set to copy-on-write;

characterized in,

the system comprises further a checkpoint module, the checkpoint module being configured to determine the one or more source segments by determining which segments have been changed in a previous checkpoint interval,

wherein the one or more snapshot segments comprise at least the segments which have been changed in the previous checkpoint interval.

EP 3 291 103 B1

2. The system (100) of claim 1, wherein the virtual memory creation unit is implemented in the kernel of an operating system of the system.
- 5 3. The system (100) of claim 1 or 2, wherein the system is configured to, on an access to a segment which maps to a persistent storage:
- allocate one or more page frames,
 - read the segment from the persistent storage to the page frames, and
 - map the segment and a corresponding snapshot segment to the page frames.
- 10 4. The system (100) of one of the preceding claims, wherein the system further comprises a memory management unit that is configured to induce a page-fault when a source segment (312, 314) and/or a snapshot segment (332, 334) is accessed.
- 15 5. The system (100) of claim 4, wherein the virtual memory creation unit (120) is further configured to perform copy-on-write, COW, in response to the page fault, wherein the COW performing comprises:
- allocating at least one page frame,
 - copying the content of one or more memory pages referenced by the accessed snapshot segments to new
 - 20 memory pages, and
 - setting the corresponding source segments to map to the new memory pages.
- 25 6. The system (100) of one of the previous claims, wherein the system comprises a compiling unit that is configured to compile a physical execution plan of a database query, in particular an OLAP query, wherein preferably the compiled physical execution plan comprises an unmap command for a snapshot of a segment at a position in the physical execution plan when the snapshot of the segment is no longer accessed during a remaining execution time of the physical execution plan.
- 30 7. The system (100) of one of the previous claims, wherein the virtual memory creation unit is further configured to provide a single page table that corresponds to the process address space, wherein the page table comprise one or more page table entries for the source segments and one or more page table entries for the snapshot segments, and the page table entries of the snapshot segments are mapped to same physical memory addresses as the page table entries of the source segments to be snapshot.
- 35 8. A method (200) for creating a snapshot of a subset of a database that is stored as a plurality of segments in a process address space, the method comprising:
- determining (210) from the plurality of segments one or more source segments (312, 314; 531-534) of the snapshot, and
 - 40 - allocating (220) for the snapshot one or more snapshot segments (332, 334) within the process address space, wherein the snapshot segments (332, 334) are mapped to same physical addresses as the source segments and wherein the source segments (312, 314) and/or the snapshot segments (332, 334) are set to copy-on-write; wherein determining the one or more source segments comprises a step of determining which segments have been changed in a previous checkpoint interval,
- 45 wherein the one or more snapshot segments comprise at least the segments which have been changed in the previous checkpoint interval.
- 50 9. The method (200) of claim 8, further comprising, on an access to a segment which maps to a persistent storage:
- allocating one or more page frames,
 - reading the segment from the persistent storage to the page frames, and
 - mapping the segment and the corresponding snapshot segment to the page frames.
- 55 10. The method (200) of one of claims 8 to 9, wherein the method is a method for maintaining a hybrid online transaction processing, OLTP, and online analytical processing, OLAP, database system and wherein the method further comprises receiving and executing an OLAP query.

11. The method (200) of claim 10, wherein determining the one or more source segments of the snapshot comprises a step of creating a physical execution plan, wherein in particular the source segments are referenced in the physical execution plan.

5 12. The method (200) of one of claims 8 to 11, further comprising:

- determining when a snapshot segment is no longer accessed by a consumer thread, and
- unmapping the snapshot segment.

10 13. The method (200) of one of claims 8 to 12, further comprising:

- setting a read-only flag of the source segments and/or the snapshot segments,
- inducing a page-fault when one of the source segments is accessed and/or when one of the snapshot segments is accessed, and
- 15 - performing copy-on-write in response to the page fault.

14. A computer-readable storage medium storing program code, the program code comprising instructions which when executed by a processor carry out the method of one of claims 8 to 13.

20

Patentansprüche

1. System (100) zur Erzeugung einer Momentaufnahme einer Teilmenge einer Datenbank, die als eine Mehrzahl von Segmenten (312, 314; 531-534) in einem Prozessadressraum gespeichert ist, das System umfassend:

25

- eine Bestimmungseinheit (110), die dafür eingerichtet ist, aus der Mehrzahl von Segmenten ein oder mehrere Quellsegmente (312, 314; 531-534) der Momentaufnahme zu bestimmen, und
- eine Einheit zur Erzeugung virtuellen Speichers (120), die dafür eingerichtet ist, für die Momentaufnahme ein oder mehrere Momentaufnahme-segmente (332, 334) innerhalb des Prozessadressraums zuzuweisen, wobei die Momentaufnahme-segmente auf gleiche physische Speicheradressen wie die Quellsegmente abgebildet sind, und wobei die Quellsegmente und/oder die Momentaufnahme-segmente auf "Copy-on-Write" gestellt sind;

30

dadurch gekennzeichnet,

das System ferner ein Prüfpunktmodul umfasst, wobei das Prüfpunktmodul dafür eingerichtet ist, das eine oder die mehreren Quellsegmente durch Bestimmen, welche Segmente in einem vorigen Prüfpunktintervall geändert wurden, zu bestimmen, wobei das eine oder die mehreren Momentaufnahme-segmente zumindest die Segmente umfassen, die in dem vorigen Prüfpunktintervall geändert wurden.

35

2. System (100) nach Anspruch 1, wobei die Einheit zur Erzeugung virtuellen Speichers im Kernel eines Betriebssystems des Systems implementiert ist.

40

3. System (100) nach Anspruch 1 oder 2, wobei das System dafür eingerichtet ist, bei einem Zugriff auf ein Segment, das auf einen dauerhaften Speicher abbildet:

45

- einen oder mehrere Seitenrahmen zuzuweisen,
- das Segment aus dem dauerhaften Speicher in die Seitenrahmen zu lesen, und
- das Segment und ein entsprechendes Momentaufnahme-segment auf die Seitenrahmen abzubilden.

4. System (100) nach einem der vorhergehenden Ansprüche, wobei das System ferner eine Speicherverwaltungseinheit umfasst, die dafür eingerichtet ist, einen Seitenfehler hervorzurufen, wenn auf ein Quellsegment (312, 314) und/oder ein Momentaufnahme-segment (332, 334) zugegriffen wird.

50

5. System (100) nach Anspruch 4, wobei die Einheit zur Erzeugung virtuellen Speichers (120) ferner dafür eingerichtet ist, "Copy-on-Write", COW, in Reaktion auf den Seitenfehler durchzuführen, wobei das Durchführen von COW Folgendes umfasst:

55

- Zuweisen mindestens eines Seitenrahmens,
- Kopieren des Inhalts einer oder mehrerer Speicherseiten, die von den Momentaufnahme-segmenten, auf wel-

che zugegriffen wird, referenziert werden, in neue Speicherseiten, und Einstellen der entsprechenden Quellsegmente, um auf die neuen Speicherseiten abzubilden.

- 5
6. System (100) nach einem der vorhergehenden Ansprüche, wobei das System eine Kompilierungs-
einheit umfasst, die dafür eingerichtet ist, einen physischen Ausführungsplan einer Datenbankabfrage, insbesondere einer OLAP-
Abfrage, zu kompilieren, wobei vorzugsweise der kompilierte physische Ausführungsplan einen "Unmap"-Befehl
für eine Momentaufnahme eines Segments an einer Position im physischen Ausführungsplan umfasst, wenn wäh-
rend einer verbleibenden Ausführungszeit des physischen Ausführungsplans nicht mehr auf die Momentaufnahme
des Segments zugegriffen wird.
- 10
7. System (100) nach einem der vorhergehenden Ansprüche, wobei die Einheit zur Erzeugung virtuellen Speichers
ferner dafür eingerichtet ist, eine einzelne Seitentabelle bereitzustellen, die dem Prozessadressraum entspricht,
wobei die Seitentabelle einen oder mehrere Seitentabelleneinträge für die Quellsegmente und einen oder mehrere
Seitentabelleneinträge für die Momentaufnahmensegmente umfasst und die Seitentabelleneinträge der Moment-
aufnahmensegmente auf gleiche physische Speicheradressen wie die Seitentabelleneinträge der Quellsegmente
abgebildet sind, die Momentaufnahme sein sollen.
- 15
8. Verfahren (200) zur Erzeugung einer Momentaufnahme einer Teilmenge einer Datenbank, die als eine Mehrzahl
von Segmenten in einem Prozessadressraum gespeichert ist, das Verfahren umfassend:
- 20
- Bestimmen (210) eines oder mehrerer Quellsegmente (312, 314; 531-534) der Momentaufnahme aus der
Mehrzahl von Segmenten, und
 - Zuweisen (220) eines oder mehrerer Momentaufnahmensegmente (332, 334) innerhalb des Prozessadress-
raums für die Momentaufnahme, wobei die Momentaufnahmensegmente (332, 334) auf gleiche physische
Adressen wie die Quellsegmente abgebildet sind, und wobei die Quellsegmente (312, 314) und/oder die Mo-
mentaufnahmensegmente (332, 334) auf "Copy-on-Write" gestellt sind;
- 25
- wobei Bestimmen des einen oder der mehreren Quellsegmente einen Schritt des Bestimmens, welche Segmente
in einem vorigen Prüfpunktintervall geändert wurden, umfasst, wobei das eine oder die mehreren Momentaufnah-
mensegmente zumindest die Segmente umfassen, die in dem vorigen Prüfpunktintervall geändert wurden.
- 30
9. Verfahren (200) nach Anspruch 8, ferner umfassend, bei einem Zugriff auf ein Segment, das auf einen dauerhaften
Speicher abbildet:
- 35
- Zuweisen eines oder mehrerer Seitenrahmen,
 - Lesen des Segments aus dem dauerhaften Speicher in die Seitenrahmen, und
 - Abbilden des Segments und des entsprechenden Momentaufnahmensegments auf die Seitenrahmen.
- 40
10. Verfahren (200) nach einem der Ansprüche 8 bis 9, wobei das Verfahren ein Verfahren zur Verwaltung eines hybriden
"Online-Transaction-Processing-", OLTP-, und "Online-Analytical-Processing-", OLAP-, Datenbanksystems ist, und
wobei das Verfahren ferner Empfangen und Ausführen einer OLAP-Abfrage umfasst.
- 45
11. Verfahren (200) nach Anspruch 10, wobei Bestimmen des einen oder der mehreren Quellsegmente der Moment-
aufnahme einen Schritt des Erzeugens eines physischen Ausführungsplans umfasst, wobei insbesondere die Quell-
segmente in dem physischen Ausführungsplan referenziert sind.
- 50
12. Verfahren (200) nach einem der Ansprüche 8 bis 11, ferner umfassend:
- Bestimmen, wann nicht mehr durch einen Verbraucher-"Thread" auf ein Momentaufnahmensegment zuge-
griffen wird, und
 - Aufheben der Abbildung des Momentaufnahmensegments.
- 55
13. Verfahren (200) nach einem der Ansprüche 8 bis 12, ferner umfassend:
- Setzen eines Nur-Lesen-"Flags" der Quellsegmente und/oder der Momentaufnahmensegmente,
 - Hervorrufen eines Seitenfehlers, wenn auf eins der Quellsegmente zugegriffen wird, und/oder wenn auf eins
der Momentaufnahmensegmente zugegriffen wird, und
 - Durchführen von "Copy-on-Write" in Reaktion auf den Seitenfehler.

14. Computerlesbares Speichermedium, das Programmcode speichert, der Programmcode umfassend Anweisungen, die bei Ausführung durch einen Prozessor das Verfahren nach einem der Ansprüche 8 bis 13 durchführen.

5 **Revendications**

1. Système (100) pour créer une capture instantanée d'un sous-ensemble d'une base de données qui est stockée sous la forme d'une pluralité de segments (312, 314 ; 531 à 534) dans un espace d'adressage de processus, le système comprenant :

- 10
15
- une unité de détermination (110), configurée pour déterminer, à partir de la pluralité de segments, un ou plusieurs segments source (312, 314 ; 531 à 534) de la capture instantanée, et
- une unité de création d'une mémoire virtuelle (120) configurée pour allouer à la capture instantanée un ou plusieurs segments de capture instantanée (332, 334) dans l'espace d'adressage du processus, dans lequel les segments de capture instantanée sont mis en correspondance avec les mêmes adresses de mémoire physique que les segments source et dans lequel les segments source et/ou les segments de capture instantanée sont configurés pour une copie sur écriture ;

caractérisé en ce que

le système comprend en outre un module de point de contrôle, le module de point de contrôle étant configuré pour déterminer le ou les segments source en déterminant quels segments ont été modifiés dans un intervalle de point de contrôle précédent, dans lequel le ou les segments de capture instantanée comprennent au moins les segments qui ont été modifiés dans l'intervalle de point de contrôle précédent.

2. Système (100) selon la revendication 1, dans lequel l'unité de création de mémoire virtuelle est réalisée dans le noyau d'un système d'exploitation du système.

3. Système (100) selon les revendications 1 ou 2, dans lequel le système est configuré pour, sur un accès à un segment qui met en correspondance avec une mémoire persistante :

- 30
- allouer une ou plusieurs trames de page,
- lire le segment de la mémoire persistante dans les trames de page, et
- mettre en correspondance avec les trames de page le segment et un segment de capture instantanée correspondant.

4. Système (100) selon l'une des revendications précédentes, dans lequel le système comprend en outre une unité de gestion de la mémoire qui est configurée pour induire un défaut de page lorsqu'on accède à un segment source (312, 314) et/ou à un segment de capture instantanée (332, 334).

5. Système (100) selon la revendication 4, dans lequel l'unité de création de mémoire virtuelle (120) est en outre configurée pour effectuer une copie sur écriture, COW (copy-on-write), en réponse au défaut de page, dans lequel la réalisation de la COW comprend les étapes consistant à :

- allouer au moins une trame de page,
copier le contenu d'une ou de plusieurs pages de mémoire référencées par les segments de capture instantanée consultés sur de nouvelles pages de mémoire, et
définir les segments source correspondants à faire correspondre avec les nouvelles pages de mémoire.

6. Système (100) selon l'une des revendications précédentes, dans lequel le système comprend une unité de compilation qui est configurée pour compiler un plan d'exécution physique d'une demande de base de données, en particulier d'une demande OLAP, dans lequel, de préférence, le plan d'exécution physique compilé comprend une commande d'interruption de la correspondance d'une capture instantanée d'un segment au niveau d'une position dans le plan d'exécution physique lorsque la capture instantanée du segment n'est plus accessible pendant un temps d'exécution résiduel du plan d'exécution physique.

7. Système (100) selon l'une des revendications précédentes, dans lequel l'unité de création de mémoire virtuelle est en outre configurée pour fournir une table de page unique qui correspond à l'espace d'adressage du processus, dans lequel la table de pages comprend une ou plusieurs entrées de table de pages pour les segments source et une ou plusieurs entrées de table de pages pour les segments de capture instantanée, et les entrées de la table

EP 3 291 103 B1

des pages des segments de capture instantanée sont mises en correspondance avec les mêmes adresses de la mémoire physique que les entrées de la table des pages des segments source à capturer de façon instantanée.

- 5 8. Procédé (200) pour créer une capture instantanée d'un sous-ensemble d'une base de données qui est stockée sous la forme d'une pluralité de segments dans un espace d'adressage de processus, le procédé comprenant les étapes consistant à :
- déterminer (210), à partir de la pluralité de segments, un ou plusieurs segments source (312, 314 ; 531 à 534) de la capture instantanée, et
 - 10 - allouer (220) à la capture instantanée un ou plusieurs segments de la capture instantanée (332, 334) dans l'espace d'adressage du processus, dans lequel les segments de la capture instantanée (332, 334) sont mis en correspondance avec les mêmes adresses physiques que les segments source et dans lequel les segments source (312, 314) et/ou les segments de la capture instantanée (332, 334) sont configurés pour une copie sur écriture ;
 - 15 dans lequel la détermination du ou des segments de source comprend une étape consistant à déterminer quels segments ont été modifiés dans un intervalle du point de contrôle précédent, dans lequel le ou les segments de capture instantanée comprennent au moins les segments qui ont été modifiés dans l'intervalle de point de contrôle précédent.
- 20 9. Procédé (200) selon la revendication 8, comprenant en outre les étapes consistant à, sur un accès à un segment qui met en correspondance avec une mémoire persistante :
- allouer une ou plusieurs trames de page,
 - lire le segment de la mémoire persistante dans les trames de page, et
 - 25 - mettre en correspondance le segment et le segment de la capture instantanée correspondant avec les trames de page.
- 30 10. Procédé (200) selon l'une des revendications 8 à 9, dans lequel la méthode est une méthode pour maintenir un système de base de données avec un traitement hybride des transactions en ligne, OLTP (online transaction processing), et un traitement analytique en ligne, OLAP (online analytical processing) et dans lequel le procédé comprend en outre la réception et l'exécution d'une requête OLAP.
- 35 11. Procédé (200) selon la revendication 10, dans lequel la détermination du ou des segments source de la capture instantanée comprend une étape de création d'un plan d'exécution physique, dans lequel en particulier les segments source sont référencés dans le plan d'exécution physique.
- 40 12. Procédé (200) selon l'une des revendications 8 à 11, comprenant en outre les étapes consistant à :
- déterminer à quel moment un filon consommateur n'a plus accès à un segment de capture instantanée, et
 - interrompre la correspondance du segment de capture instantanée.
- 45 13. Procédé (200) selon l'une des revendications 8 à 12, comprenant en outre les étapes consistant à :
- définir un indicateur en lecture seule des segments source et/ou des segments de capture instantanée,
 - induire un défaut de page lors de l'accès à l'un des segments source et/ou lors de l'accès à l'un des segments de capture instantanée, et
 - effectuer une copie sur écriture en réponse au défaut de page.
- 50 14. Support de mémoire lisible par ordinateur stockant du code de programme, le code de programme comprenant des instructions qui, lorsqu'elles sont exécutées par un processeur, mettent en œuvre le procédé selon l'une des revendications 8 à 13.
- 55

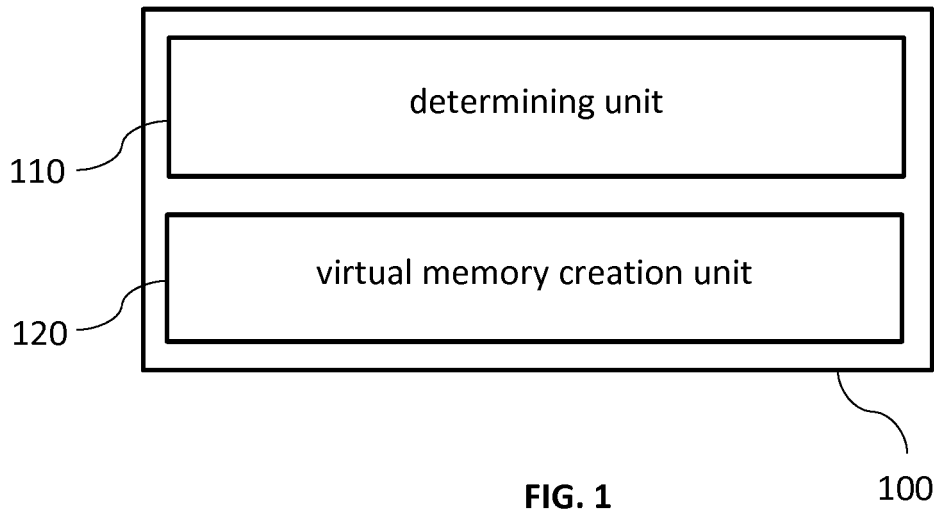


FIG. 1

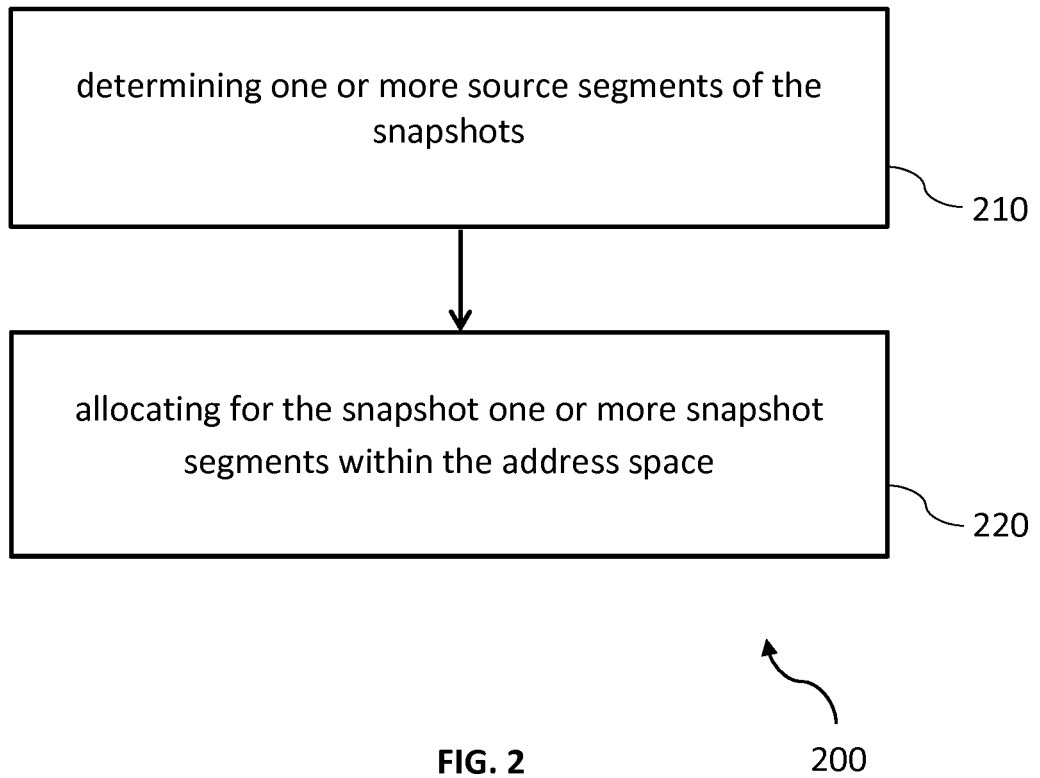


FIG. 2

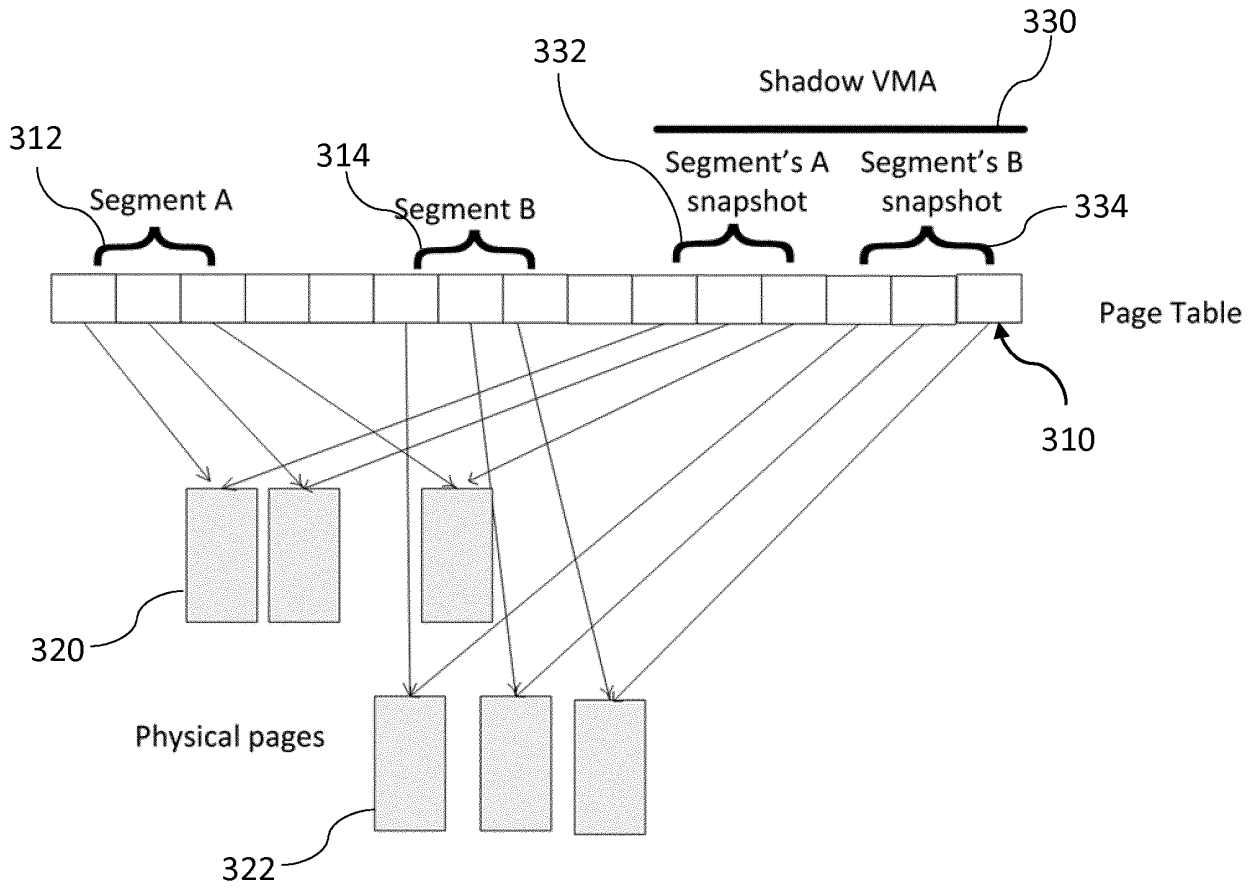


FIG. 3

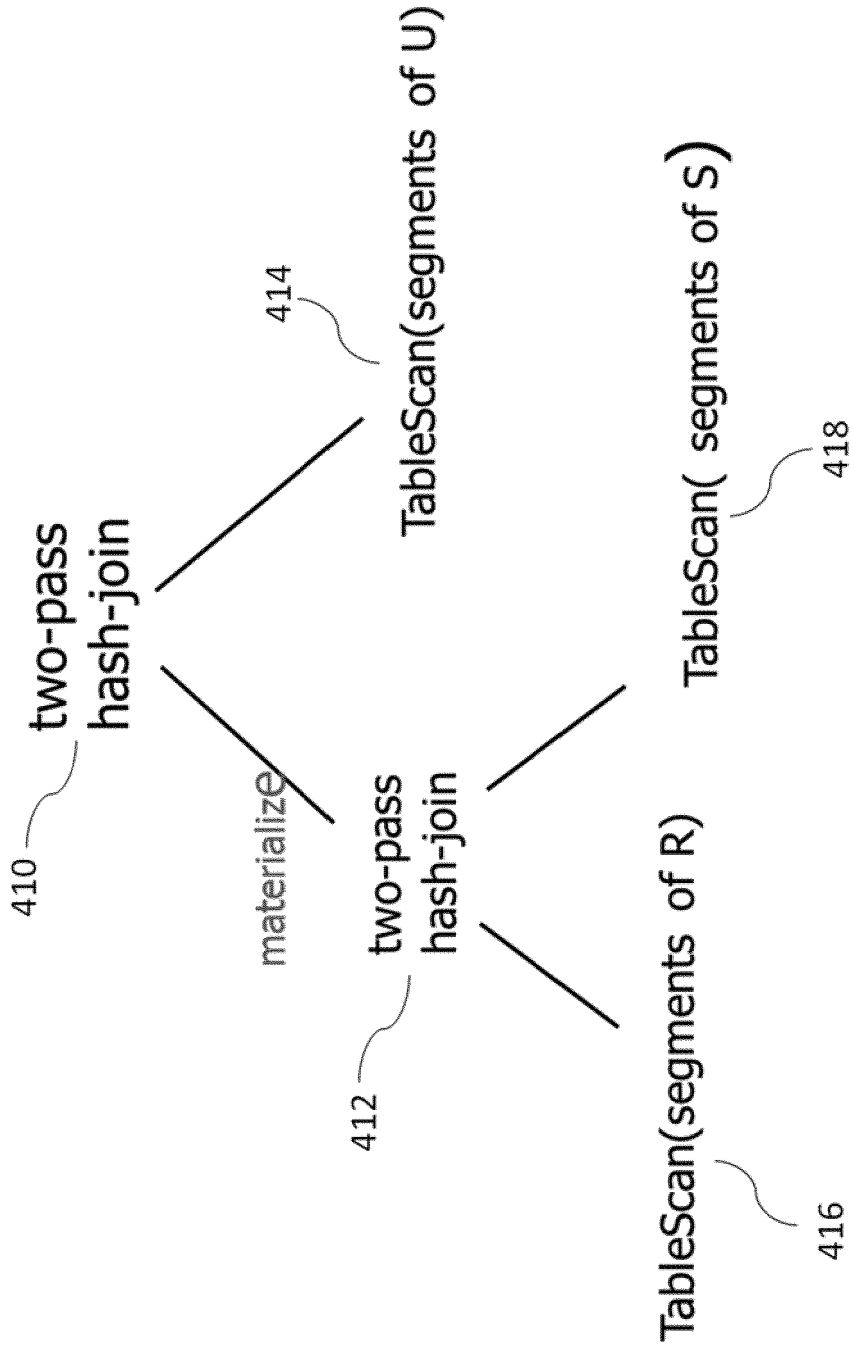


FIG. 4

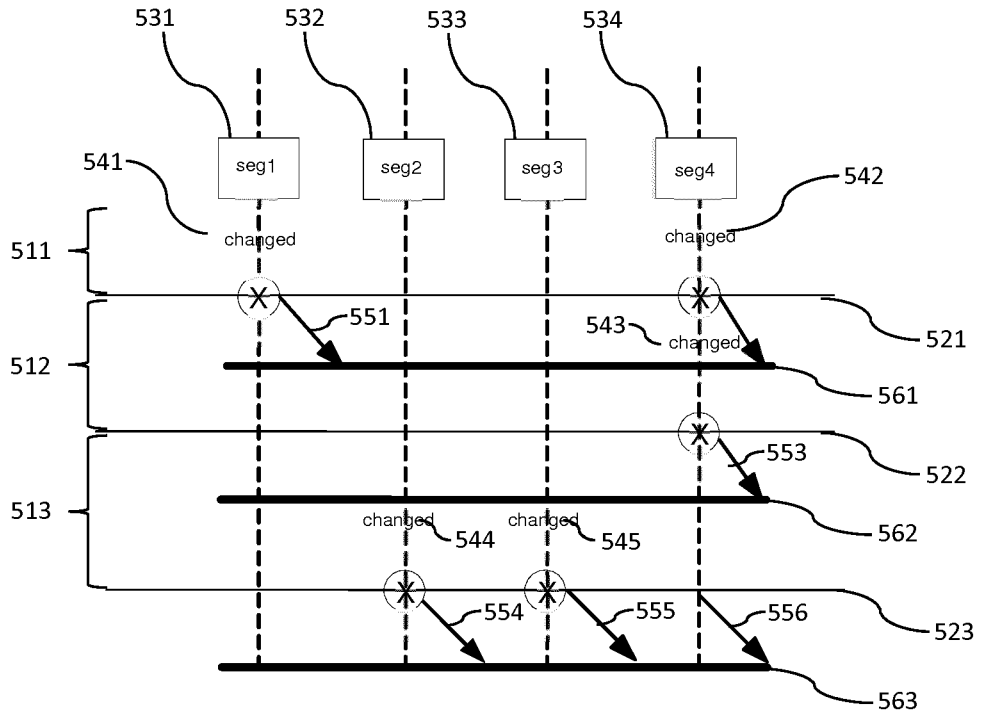


FIG. 5

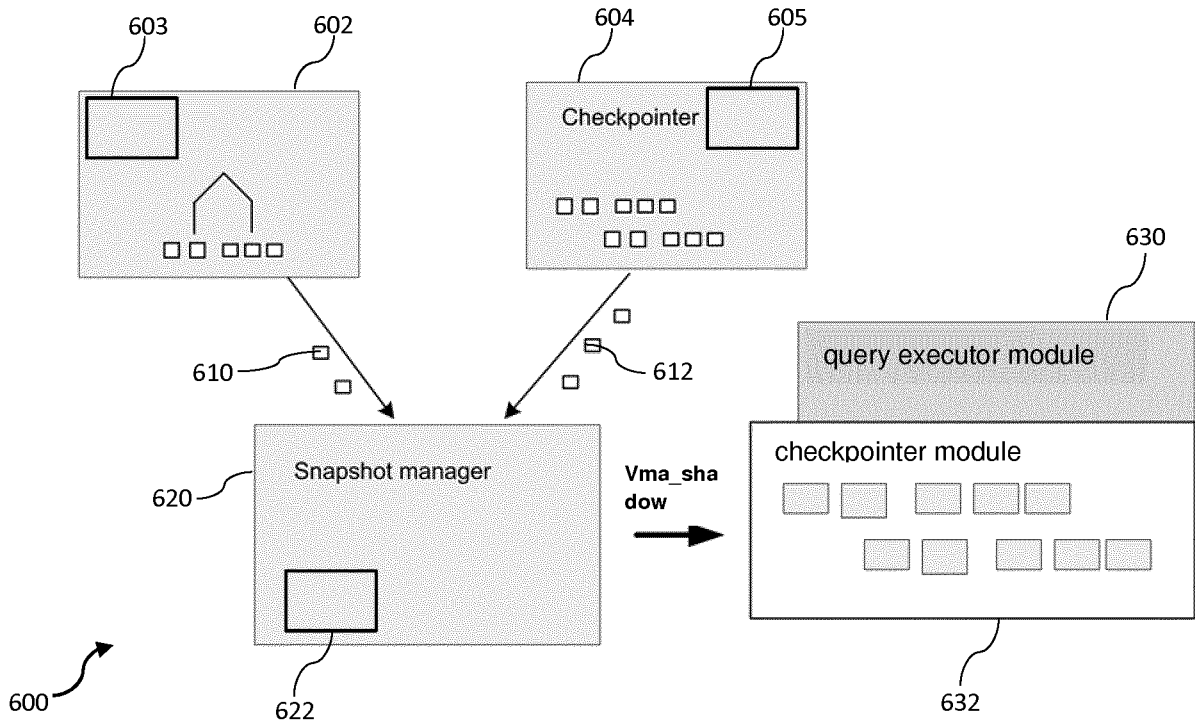


FIG. 6

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 2013073513 A1 [0007]

Non-patent literature cited in the description

- Supporting Persistent Object Systems in a Single Address Space. *Proceedings of the 7th International Workshop on Persistent Object Systems*, 1997 [0006]